



Innovative Computing Laboratory
UNIVERSITY OF TENNESSEE
COMPUTER SCIENCE DEPARTMENT

Metacomputing: What do Users Want?

Dr. David Cronk
Innovative Computing Lab
University of Tennessee
June 13, 2002

Outline


Introduction to Metacomputing

What types of users are there?

What do different types of users want?

Conclusions

Metacomputing



Metacomputing (cont)

Allow large-scale applications to make use of a collection of HPC resources in a seamless manner

A single problem may execute on multiple supercomputers simultaneously

Practically speaking, spanning resource is impractical at this time

Metacomputing (cont)

Systems such as Legion and Globus are large, monolithic systems that do little to hide the grid's complexity

Many simpler to use systems are layered on top of these large systems

These simpler systems offer a variety of services

Caution! Interfaces change!

Network Enabled Servers (NES)

Provide abstractions for building grid applications

Provide access to both hardware and software resources

Provide this access through RPC type services

Libraries can be accessed, with no worry about compilation, maintenance, etc.

Obstacles for Grid Systems

Binary compatibility

Users do not want to maintain binaries for lots of different machines

File staging and data management

If you do not know where a job will run, you cannot pre-stage the data

Results should be readily available to user

Scheduling

What is the primary goal?

Users

Unsophisticated Users

Does not refer to personal sophistication, but computer savvy

Intermediate Users

Know how to use HPC resources

Sophisticated Users

Have deep understanding of many hardware and software issues

Users (cont)

Unsophisticated Users

See their computers mostly as black boxes used to generate results to be analyzed

Do little or no programming, often working with codes written and maintained by others

Make heavy use of numerical libraries and interactive environments

Do little or no algorithm development

Users (cont)

Intermediate Users

- Do some programming and algorithm development

- Understand how their codes work and perform some optimizations

- Does a lot of HPC, but generally works with codes parallelized by others

- Understands how to submit jobs to HPC resources and how to gather results

Users (cont)

Sophisticated users

- Embrace new technologies

- Put forth significant effort to improve the performance of their code

- Often hand parallelize sequential code and/or optimize parallel algorithms

- Have intimate understanding of how their code works

- Often optimize codes for specific architectures

What do Users Want?

Often, not metacomputing!

Many users are perfectly happy with the status quo

Unsophisticated Users

Little or no interest in metacomputing

Might be interested if they could run larger problem sizes or generate results more quickly

Would need to have simple access to results

This would all need to be done with a **VERY** small learning curve

What do Users Want?

Intermediate Users

These users tend to be willing to put forth moderate effort to learn a new system

Actual use of the system after the learning curve should be little or no more difficult than current methods

What do Users Want?

Sophisticated Users

These users tend to be willing to put forth significant effort to learn a new system

Willing to deal with more complicated job submission protocols once the learning curve has been completed

What do Users Want?

Intermediate and Sophisticated Users

- Not interested in exploring metacomputing just because it is new and interesting

- Still most interested in their science

- Only interested if real, tangible benefits are available

- Must improve productivity

What do Users Want?

Intermediate and Sophisticated Users

Improved completion time

- Don't really care how long the job runs

- Time from job submission to when results are available is what's important

Easier file staging

- These users are often using very large data sets

- Ensuring data is available become tedious and time consuming

Conclusions

There is very little interest in metacomputing among the end-user community

This lack of interest is most pronounced among unsophisticated users

These users are only interested in systems that are VERY easy to learn

NES or similar systems that require little code alteration may be best for these users

Conclusion (cont)

Intermediate and sophisticated users are already using HPC resources, meaning NES and similar systems offer little in terms of performance

Systems that offer improved scheduling may lead to faster turn around time

Seems unlikely any system will offer actual improved runtime in near future

Systems that make file staging easier are very attractive

Conclusion (cont)

Many users are satisfied with the status quo

Their current methods may not be the “best”, but they are “good enough”

Some users (of all levels) have simply stated there is nothing a grid system could offer to make them use it

Conclusion (cont)

It seems that we are in a catch-22

Users will not be interested in metacomputing until there are systems that meet their needs well enough to warrant the interest

We cannot design systems that meet users' needs unless they have enough interest to let us know what those needs are

Conclusion (cont)

Blindly guessing what users want, implementing this, and handing it to the user rarely works

If metacomputing is going to be successful within the DoD, effort needs to be made to increase user interest

More interest will allow the DoD to better define the users' requirements

This work serves as a starting point, but due to general apathy towards grid computing, it is difficult to gather comprehensive results